UCRL-TR-200982

On the Optimal Allocation of R&D Resources for Climate Change Technology Development

G.J. Blanford, L.E. Clarke

November 17, 2003

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

On the Optimal Allocation of R&D Resources for Climate Change Technology Development

Geoffrey J. Blanford
Department of Management Science and Engineering, Stanford University
and Lawrence Livermore National Laboratory

Leon E. Clarke
Lawrence Livermore National Laboratory

November, 2003

Abstract

While technology studies and integrated assessment models incorporating endogenous technological change have demonstrated that advancing technology is a crucial component of an optimal greenhouse gas abatement strategy, the R&D process itself has received little analytical attention. This paper presents a conceptual framework for considering and exploring the optimal allocation of R&D resources for climate change technology development. The framework is then applied to a stylized application that considers the allocation between R&D focused on resolving uncertainties about the retention of sequestration and R&D focused on improving the performance of renewable energy technologies.

1 Introduction

Technological change is a cornerstone of virtually every interpretation of an appropriate response to the challenge of global climate change. Even with lenient stabilization targets and curbed growth scenarios, massive deployment of energy sources with no greenhouse gas (GHG) emissions will be required over the coming century, on the order of today's entire energy system (Hoffert et al. 1998, Caldeira et al. 2003). Many believe that the current resource base and state of technical knowledge are unable to support deployment of this magnitude at a reasonable cost (Hoffert et al. 2002). While socioeconomic measures may be able to effect demand reductions and altered consumption patterns, and while significant emissions abatement can be cost-effectively achieved with known technological options (IPCC 2001), the development and introduction of new technology is distinct from these methods in its far-reaching ability to facilitate an affordable long-term transition.[1]

Given technology's crucial role in achieving stated climate policy goals in the U.S. and internationally,[2] it is natural to examine carefully the implications of a strategy for its development. In particular, although privately funded research and its interactions with public policy are fundamental to technology development, it is useful to approach the problem from the "social perspective". This is usually interpreted as that of the government, but may also be viewed as the examination of global resource allocation regardless of funding source and incentives. A public research and development (R&D) strategy can be roughly divided into two concerns: (1) the range and extent of research to be undertaken, and (2) the public agency role in facilitating this R&D. The second concern involves not only interactions between public and private R&D, but also interactions between public agencies (e.g., California and the Department of Energy), and interactions between countries. The focus of this paper is on the first concern: in the broad social sense

[1]Technological advance is essential even from the skeptic's perspective—arguments for maintaining the current suite of energy technologies often depend heavily on the belief that human ingenuity will succeed in producing adequate adaptive measures in the future should any damages be realized.

[2]The United Nations Framework Convention on Climate Change (UNFCCC) has resolved to stabilize greenhouse gas concentrations "at a level that would prevent dangerous anthropogenic interference with the climate system" (1992; see http://unfccc.int/resource/docs/convkp/conveng.pdf). President Bush affirmed the United States' commitment to this goal in speeches on June 11, 2001 and February 14, 2002.

described above, determining the optimal allocation of R&D resources.

Within this context, there are many important avenues of inquiry. For example, two technology areas expected to figure prominently in coming decades are renewable energy (primarily wind and photovoltaics) and carbon capture and sequestration. This paper seeks to address strategic R&D questions such as: Should development of renewable energy technologies be postponed until the retention rates of geologically sequestered carbon are better understood? Should near-term investment in resolving this uncertainty increase in the belief that sequestration will have high retention rates or the effectiveness with which R&D can discover this? Should the portfolio be adjusted more toward renewable energy or more toward sequestration in response to increasing uncertainty about the appropriate stabilization target?[3]

We recognize that explicit answers to this type of inquiry not only require representation of highly complex systems, but also are inherently based on subjective judgements of uncertainty and in some cases value. However, even with a simplified representation and some subjectivity, we propose in this paper that it is possible to better understand the forces that drive the answers. Previous work applying integrated assessment models (IAMs) to the climate problem has provided critical insight into the *value* of technological advance, but these models do not capture the R&D process itself—e.g., the productivity of R&D investment, the allocation of resources under uncertainty about research success, or the decision between different types of research. We believe that the climate change R&D planning problem now warrants the same level of rigorous treatment that has been applied to other aspects of climate policy, such as optimal emissions trajectories and optimal carbon taxes.

This paper has three primary objectives. The first is demonstrate the need for formal research on R&D planning in the climate context that supplements, complements, and interacts with the extensive work in integrated assessment modeling. The second objective is to present a relatively simple conceptual framework to meet this need. The framework we present puts heavy emphasis on R&D decisions, while treating crudely the economic interactions between technologies, typically the focus of IAMs. The third objective of the paper is to show that interesting and important insights are available through more explicit consideration of the R&D planning problem. To this end, we explore

[3]Uncertainty about the appropriate stabilization target is used here as a proxy for uncertainty about the totality of impacts from continued GHG emissions.

a stylized application that considers the allocation between R&D focused on resolving uncertainties about the retention of sequestration and R&D focused on improving the performance of renewable energy technologies.

The conceptual framework we employ is designed to inform and eventually incorporate the results of IAMs. However, the R&D process model component is not explicitly linked to an IAM—integration is not necessary in our framework. The goal of this study is not to solve the R&D planning problem, but rather to bring to it greater rigor, so that the strengths of IAMs and other studies of the climate problem may be brought to bear more directly on the R&D policy process. The structure we propose is general, but will at a minimum enable (1) an understanding of the connection between R&D investment decisions and technological advance; (2) collection and utilization of information on technological potentials, barriers, and limits along with empirical observations of the productivity of R&D itself; and (3) identification of key future scenarios and uncertainties to be explored through IAMs or other approaches.

The paper proceeds as follows. In Section 2, we discuss previous research relevant to R&D portfolio planning for climate change technologies. In Section 3 and Section 4, we discuss two background issues that play into the R&D portfolio planning process, namely the various roles of R&D and the reasons for pursuing a diversified portfolio. Section 5 introduces the our conceptual model. Section 6 introduces a stylistic application of the conceptual model that focuses on an allocation between renewables, and sequestration. Section 7 presents the results of the stylistic application; Section 8 concludes.

2 Previous Research

To date, although there is a great deal of research relevant to both the R&D portfolio problem and the climate change context, there has been little research directly on the problem in this particular context with formal or stylistic models. Here we discuss several areas of research that can support the climate change technology R&D planning problem.

An important backdrop to any problem in the climate context is the extensive analysis conducted with IAMs.[4] These models explore the long-term implications of trajectories of technological advance and the interac-

[4]For example, see Manne et al. (1993), Nordhaus (1994), Edmonds et al. (1994), and Prinn et al. (1999).

tions between technological advance and optimal climate policy or emissions trajectories. IAMs are ideally suited to understanding the value, both in terms of costs and avoided emissions, of particular technological scenarios because of their comprehensive representation of the physical and economic systems in which technologies operate and compete. However, until recently technological advance has entered exogenously into most IAMs, so that they are unable to consider the underlying R&D process. From an *ex poste* perspective, integrated assessment analysis frequently does indicate that a large variety of GHG-free technologies might contribute to stabilization. However, it is misleading to make inferences from these predictions about the optimal diversification of the current research portfolio. This issue is discussed in Section 4.

In recent years, with the realization of the importance of technological advance for climate stabilization, there have been increasing efforts to incorporate endogenous technical change (ETC) into IAMs.[5] In these models, the aggregate investment in innovation is a decision variable, capturing the response of private actors to policy incentives, price signals, and learning benefits from deployment. These contributions imply that substantial innovation investment can lower the total cost of stabilization, and that optimal near-term emissions policy may be less stringent when accounting for increased innovation.[6] However, in their current formulation, ETC models do not consider allocation of innovation resources across technologies, and represent returns to investment as deterministic, whereas our interest here is in characterizing the optimal R&D portfolio when returns are uncertain. Still, our study relies on the ETC literature for the result that *some* R&D in-

[5]Efforts to endogenize technological change generally use one of two approaches—R&D based advance and learning-by-doing. Important examples of the former include Goulder and Mathai (2000), Schneider and Goulder (1997), Goulder and Schneider (1999), Nordhaus (2002), and Popp (2002). Important examples of the latter include Seebregts et al. (1999), Gritsevskii and Nakicenovic (2000), Grubler and Gritsevskii (2001), Goulder and Mathai (2000), and van der Zwaan et al. (2002). Also, see Romer (1990) for a theoretical framing of ETC.

[6]There is significant debate over the impact of ETC on the near-term stringency of climate constraints. Models that include R&D as a decision variable along with emissions levels indicate that the more innovation that is feasible, the less stringent should be near-term actions, so as to take advantage of later technological advances. However, models based on learning-by-doing indicate no clear direction (see Goulder and Mathai 2000). There has been no rigorous consideration of this question when R&D is imperfectly stimulated by emissions policy.

vestment is optimal, since we focus on relative rather than absolute resource allocation.

The economics literature provides another angle of support with its extensive coverage of portfolio theory and its theoretical research on private-sector innovation, in both environmental and non-environmental technologies. While most of this literature again focuses on the R&D behavior of private actors,[7] there have been several efforts explicitly considering portfolio concerns in this context. For example, Dasgupta and Maskin (1987), in exploring optimal diversification of R&D portfolios, implicitly demonstrate that an assumption, widely made in this literature, of decreasing near-term returns to innovation investment results in a more diversified optimal portfolio. Loch and Kavadias (2003) illustrate similar effects of this assumption, although they argue that it is most applicable to private firms operating in mature markets. Section 4 further explores this topic.

One piece of research that explicitly considers R&D portfolio issues in the climate context is Baker et al. (2003). The authors incorporate a simple, one-period model of R&D into a stochastic version of the DICE model (Nordhaus 1994) to ascertain the impact of increasing damage uncertainty on optimal near-term R&D. The results indicate that R&D into high-cost, very-low emissions technologies such as photovoltaic (PV) cells can serve as a hedge against bad climate outcomes, so that optimal investment in such technologies will likely increase in uncertainty about damages.

3 The Objectives of R&D

R&D serves several objectives, some of which may not be explicitly intended to advance technology. Here we discuss three of these objectives. We believe that analysis of an R&D portfolio for climate change should consider variation in R&D purposes in addition to variation across technology area. Moreover, it is important to the design of a planning model to understand the different ways in which R&D can provide benefits.

The most obvious role for R&D is to advance technology. For example, research on PV cells is focused on reducing their costs or increasing their efficiency at converting sunlight to electricity. Within this broad objective

[7] As mentioned above, while private sector behavior is important to R&D planning, this paper considers the full level of R&D, including both public and private. It does not consider the allocation between these sources.

exists a well-known spectrum. At one end is relatively deterministic applied research, or development, to bring concepts from the bench to application. At the other end of the spectrum is basic research that has highly uncertain long-term impacts on technology. Despite the uncertainty, basic research has been consistently proven to carry large long-term benefits. The prototypical research program in renewable energy technologies examined here is cast in this mold.

A second motivation for R&D, or a second category of benefit, is to resolve uncertainty. Such a resolution may be the primary goal of the research program, or it may be a secondary benefit from a more conventional program. In particular, we identify three types of uncertainties R&D may address, directly or indirectly: (1) uncertainty about the potential for technological advance; (2) uncertainty about exogenous factors affecting the value of technological advance; and (3) uncertainty about some aspect of the technology itself.

The first type of uncertainty relates to how far a technology might ultimately advance, how much it might cost to achieve this advance, and how long it might take. When R&D intended to advance a technology is conducted and the results observed, information is obtained about the potential for advance with continued investment. For example, heavy expenditure without results over several years might suggest that the chances of future improvements would be smaller, or at least that expectations for progress by a certain deadline should be more modest. As discussed below, these ancillary incentives are partially represented in our framework. The second type, exogenous uncertainty, is abundant in the climate context. For example, there is a set of uncertainties regarding the link between emissions and human welfare, including uncertainty about climate sensitivity, the impacts of sea-level rise on human welfare, the effects on the spread of disease, and so forth. While R&D specifically directed toward understanding these phenomena is not considered in our model, the uncertainties themselves have important implications for the allocation of innovation resources to other programs.

The third type of uncertainty is concerned with aspects of current technology performance that R&D might resolve. The most significant example in the climate change context is the permanency of sequestered carbon. Planning for climate change might depend critically on the resolution of this uncertainty. In our model, a research program designed exclusively for this purpose is compared to the more traditional renewable research program.

Also important may be uncertainties about resource bases, such as such as the supply curve for uranium or the availability of geological reservoirs for carbon sequestration.

A final reason to conduct R&D is to maintain capabilities. It is crucial for firms and other R&D actors to be able to respond to new information or ideas that might emerge from other firms, other parts of a firm, other industries, or other countries. Research indicates that many firms maintain R&D investment largely to take advantage of opportunities that arise elsewhere or to be aware of these opportunities. There is empirical evidence that firms may be willing to spend more to develop a new product than to buy the same technology from a competitor, presumably because the perceived value of knowledge more than offsets the premium.[8] This purpose is not explicitly modeled here; however, it may be straightforward to capture this type of behavior by enforcing a minimum investment in each program under consideration.

4 The Basis for a Diversified Portfolio

It is frequently stated as a matter of course that diversification is a critical characteristic of an optimal climate change technology R&D portfolio. While diversification is generally a beneficial characteristic for investment portfolios, including R&D portfolios, the evidence in support of a diversified R&D portfolio specific to the climate change technology context has not, to date, been well articulated. Moreover, diversification is an inexact characteristic, defined loosely as exhibiting some measure of breadth (i.e., a wide range of projects) and balance (i.e. equalization of effort across projects), so that many possible portfolios could be classified as diversified. Hence, it is perhaps more productive to cast the problem not in terms of diversification, but in terms of the optimal allocation of R&D resources among possible research projects or technologies. The stylized application in Section 6, through extrapolation from a pairwise comparison, seeks to quantitatively investigate and begin to characterize this optimality. As background, this section presents and discusses several systematic arguments that have been,

[8]See Rosenberg (1990). In the context of a carbon policy, this phenomenon can apply to the trade-off between performing emissions reductions and purchasing permits—firms may be willing to finance reductions more costly than the equivalent permit price so as to gain the benefit of experience.

or potentially could be, raised in support of a diversified climate change technology R&D portfolio, the caveats above notwithstanding. These justifications fall into three classes: decreasing returns to scale, risk management, and heterogenous applications.

In the case of decreasing returns to scale in R&D, the marginal productivity of innovation investment is decreasing the more that is invested. Applying the standard optimality criterion of equal marginal returns across investment options, a single research program will not dominate, since, as its marginal returns fall with increased funding, it will be optimal to move funds to programs with higher marginal returns to the first dollars spent. Thus investment is spread across several programs to exploit the most productive range in each.[9] In contrast, with increasing returns and no uncertainty, marginal productivity is increasing, so that one program will eventually dominate all others and command the entire optimal allocation. These results are demonstrated in both Dasgupta and Maskin (1987) and Loch and Kavadias (2003). Especially in the near-term, the assumption of decreasing returns to scale is common throughout the economics literature on innovation, and is incorporated into the applied model presented in Section 6.

Risk management refers to the desire to decrease "risk", broadly defined, imposed by the wide variety of uncertainties associated with the climate change onto the R&D planning problem. Concepts such as "insurance" and "option value" fit into this category. The classic analog from the finance literature is the minimum variance associated with a particular expected portfolio return.[10] A large contributing factor to the risk inherent in the climate context is the uncertainty about the ultimate welfare consequences of continued GHG emissions. Optimal response to this type of risk might include diversification to the extent that different technology mixes will be most useful under different resolutions of the uncertainty.[11] Uncertain technological returns to R&D investment also expose the portfolio to risk. Diversification provides

[9]This argument does not hold in all cases. For example, technologies may receive zero funding if they have lower marginal returns at zero investment than the marginal returns from funded technologies at optimality.

[10]Note that there are other measures of risk beyond variance. See, for example, the large literature on stochastic dominance.

[11]For example, a quick transition to a renewable-heavy mix may be justified if the welfare consequences are very large, whereas a slow transition from fossil fuels might be justified if the consequences are mild. A diversified portfolio can minimize risk by attacking both of these possibilities. This topic is addressed in Baker et al. (2003).

insurance against any single technology not advancing as far as expected. Finally, maintaining a diversified portfolio of capabilities, thereby providing the ability to take advantage of opportunities in any particular technology area, can also be justified by uncertainty about when and where unexpected, exogenous technological breakthroughs might arise.

Heterogenous applications refers to the idea that, to the extent that a heterogeneous mix of technologies will be deployed in the future, irrespective of technological advance,[12] R&D funds should be correspondingly distributed across the candidate technologies. While the first two justifications for diversification have a solid logical foundation, this argument can be misleading. First, predictions of heterogeneous technology mix made by many IAMs arise from built-in random choice assumptions meant to model observed diversity, and are therefore not based on the application-specific constraints that are likely to actually produce heterogeneity. More to the point, belief that a particular technology will have significant future deployment does not necessarily mean it should receive significant current R&D funding. The optimal allocation of R&D resources depends more directly on the potential for technologies to improve, and on the effectiveness of R&D funds at achieving this improvement. Certainly it would be unwise to invest in a program that may develop quickly, but has little potential for deployment; we argue here that the reverse is also true.

Thus arguments about when diversification of innovation investment is appropriate depend fundamentally on a detailed characterization of the technologies involved. Without careful consideration of issues such as decreasing returns, productivity of innovation investment, and long-term potential for advance and deployment, even the case for a diversified portfolio lacks support, underscoring the need for more rigorous approaches to the R&D planning problem.

5 The Conceptual Framework

Figure 1 illustrates schematically the conceptual framework for the R&D planning model proposed in this paper. The framework models the invest-

[12] Heterogeneity arises because of variations in technology applications. For example, PV cells are advantageous in locations isolated from the electricity grid, whereas hydroelectric power is valuable in areas with significant rainfall and appropriate terrain. In general, it is possible that no one carbon-free energy technology can supply demand in every context.

ment allocation decision between different research programs over a finite time horizon, and has two components: (1) a sequential and uncertain R&D process model extending over a moderate time horizon (20 years in the stylized example), in which the decision variables are the levels of funding received by individual research programs in each time period; and (2) a discrete set of possible R&D outcomes that are valued externally by the expectation of a utility-weighted objective.[13] While the framework could be used to ascertain an "optimal" portfolio, the real value of exercising such a framework is to gain insights into the dynamics of investment decisions and to identify the key factors to which they are sensitive.

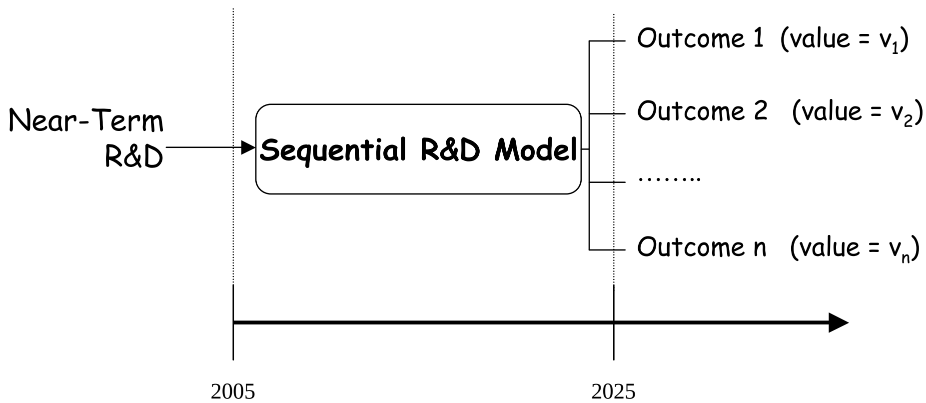

Figure 1: A Conceptual Overview of the Modeling Structure

The R&D process model links R&D investment to the resolution of technological uncertainties and to technological advance. The model assumes uncertain returns to investment, so more than one outcome is possible from

[13]Another approach would be to incorporate the R&D process model directly into an IAM. R&D decisions would be made every period along with emissions and capital investment decisions in the IAM. This explicit linkage has been used already in deterministic models such as those in Nordhaus (2002), Popp (2002), and Goulder and Schneider (1999), but without considering individual research programs. While an explicit linkage has advantages, it also has several disadvantages. For one, the explicit linkage reduces flexibility to consider a wide variety of R&D factors. For example, the detail of the R&D modeling must be consistent with the detail of the IAM in consideration, whereas our goal is to develop a framework that might be linked, at least implicitly, to a variety of IAMs. A second weakness is the difficulty detailed IAMs have had considering sequential decision-making under uncertainty, critical to the R&D planning process.

a particular path of R&D spending. By including multiple periods, the R&D process model can capture the dynamic resolution of uncertainty and the response to this resolution through changed research allocations. A wide variety of models would fit within this framework, from highly mathematical stochastic control models to classic decision analysis models. From a practical standpoint, the R&D process model should be designed so as to facilitate collecting and working with real technological data or subjective data solicited from experts.

An important characteristic of the framework is that the possible research outcomes are given a discrete rather than a continuous distribution—it is fundamentally easier to characterize, understand, and value specific, discrete outcomes than continuous distributions. The full range of possible research outcomes may be aggregated into a small number of technological scenarios, as in the application in Section 6, to simplify intuition. Each outcome is characterized by a state of technology and a state of knowledge about future potential; thus the outcomes represent intermediate scenarios in the course of a century-long stabilization effort.

The utility weighting is applied to these aggregated scenarios in order to capture a set of relative preferences among the different states of technology and knowledge. The scenarios can be valued both informally through expert solicitation and review of existing IAM scenarios, and formally through explicit linkage with one or more IAMs. The values of the scenarios may vary based on uncertain, exogenous factors, such as climate sensitivity or economic growth. We consider the creation, valuation, and ranking of these possible scenarios to be one of the more important contributions of this research approach.

Two additional consequences of the framework's design deserve note. First, while the approach allows for explicit modeling of the R&D process, its timescale is far shorter than the timescale of the climate problem, usually assumed to be at least a century. This means that further technological progress is an uncertainty, along with all the other uncertainties involved in the climate context, and can only be represented as a contributing factor the utility valuation.[14] Second, because the modeling approach separates R&D from the systems modeled by IAMs, it is not possible to capture the interactions between R&D and policy during the time that R&D is ongoing. If, for

[14]This method of incorporating uncertainty into utility valuations has been used before in practical decision analysis applications (for a discussion, see Keeney and Raiffa 1976).

example, R&D resolves in 2015 that sequestration will have very high leakage rates making it untenable as a long-term technology solution, emissions policy could not react until 2025. In the other direction, if climate science were to indicate in 2015 that there is a much higher than expected chance of extreme events or that population and emissions will naturally grow more slowly than we originally thought, the R&D portfolio cannot be adjusted in response.

6 A Stylistic Application

This section presents a stylized example of the conceptual framework. It models the allocation decision between research to resolve uncertainties about the permanency of geological sequestration and research to improve the cost and performance of renewable technologies such as wind and solar power.[15] Section 6.1 presents the R&D process model. Section 6.2 discusses the valuation of the outcomes of the process model.

6.1 The R&D Process Model

The R&D process model covers twenty years in two ten year periods. A fixed research budget each period is to be allocated between two research programs, which are different in structure and theme. Figure 2 depicts single-period trees for the two research programs. The first program aims to resolve technological uncertainty regarding the permanency of geologic carbon sequestration. The program may succeed or fail in either period, where success corresponds to an uncertainty resolution, either positive (the technology is proven viable) or negative (a fundamental flaw is demonstrated). Failure in the sequestration program therefore corresponds to a lack of resolution, rather than technological failure. We assume that if sequestration proves viable, it will be economically attractive and would therefore constitute a first-best outcome. Hence, the major concern with sequestration is permanency and not cost. For the remainder of this paper, we assume an *a priori*

[15] Obviously this omits many important climate change technologies. There are two interpretations of this abstraction. One is to view the model as a pairwise comparison of the allocation decision, only part of the larger allocation process. A second is to view the "renewables" program as the entire suite of non-fossil energy technologies, so that the model captures the full allocation decision.

subjective assessment, denoted by q in Figure 2, of 50 percent for the conditional probability that given a resolution, the sequestration uncertainty will resolve positively.[16]

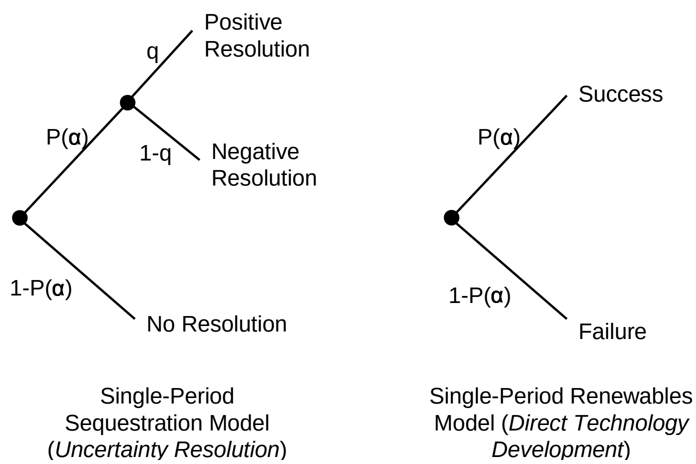

Figure 2: Sequential Decision Tree Framework

The second program aims to achieve substantial technological improvements in renewable energy technologies such as solar power. Success each period is binary, and therefore corresponds to the more traditional notion of the realization of a technological breakthrough. Failure in the renewable program corresponds to technology remaining in its current state. At this point in the model development, it is unnecessary to define the actual cost or efficiency numbers that would result from advance in renewables. We need only define the utility of the outcomes of advance. We assume that two successes in renewables will lead to a first-best situation analogous to a positive resolution of sequestration. That is, renewables will be an economically attractive approach to climate change. A single success in renewables will be less valuable, but would noticeably reduce the costs of quick climate reductions in comparison to today's technological options.

Investment affects the probability, rather than the extent, of success for each research program. The probabilities of moving down branches in the

[16]This conditional probability is the same in both periods, implying that failure to resolve the uncertainty does not influence beliefs about which direction an eventual resolution might point.

respective trees are based on the portion of the research budget allocated to each program, α_i, where i represents the research program, and $\alpha_{seq}+\alpha_{rnw} = 1$. Specifically, the probability of success in research program, P_i, is defined by the following innovation production function:

$$P_i = \rho_i \sqrt{\alpha_i} \tag{1}$$

where the parameter ρ_i represents a limiting success probability for research program i—the maximum probability of success for the program if all funds were allocated to it. The limiting parameter is an intrinsic characteristic of the research program or technology it represents. This particular functional form is entirely arbitrary, but is attractive for this prototype application because of its analytic simplicity. In this paper, we do not attempt to make a single estimate of the ρ parameters, but instead assess the sensitivity of the results to a wide range of values.

A fundamental assumption of the model is that the relationship between investment and probability of success is concave. In other words, innovation effort exhibits decreasing returns to scale, or $\frac{\partial^2 P}{\partial \alpha^2} < 0$. As discussed in Section 4, this property is a primary motivation for diversification, since an optimal allocation will equalize marginal productivity across research programs.

By employing multiple periods, the model captures the trade-offs between the two research programs. For example, the benefit of a resolved uncertainty in the first period lies in the ability to make better subsequent decisions. In this model, a negative resolution in the first period for the sequestration program leads to a second period investment entirely in the renewable program, whereas with first period failure to resolve sequestration uncertainty, the second period decision may continue to consider both programs. When only one period is considered, there is no distinction between a negative resolution and a failure, since both result in the same technological outcome. Another important implication of the sequential decision framework over a finite time period is that past successes and failures determine at any intermediate decision point the remaining potential in a particular research program; thus updated expectations inform the dynamic investment allocation decision. For example, failure in the first period in renewables implies that only one success is possible before the end of the 20-year time horizon. As mentioned in Section 3, the model captures benefits from R&D in terms of both technological advance and information about future potential.

For an intuitive interpretation of the parametrization of the model, consider the following illustrative numerical example. Assume that the budget to be allocated consists of total worldwide investment in climate change technology R&D, including both private and public funding, and that this budget amounts to \$4 billion annually.[17] The R&D decision is how to allocate these funds between advancing renewables and resolving the uncertainty about the permanency of sequestration. We illustrate three simple instances of an investment strategy which may or may not be optimal: invest entirely in sequestration, entirely in renewables, or an even balances of the two.

Suppose $\rho_{seq} = \rho_{rnw} = 0.5$. With these parameters, if all \$4 billion were spent on sequestration every year for the next twenty years, that is, if 100 percent of the allocation in both periods were committed to this program, there would be a 75 percent chance that the uncertainty would be resolved before the end of the time horizon, with a 50 percent chance that this would lead to a positive outcome and a 50 percent chance that it would lead to a negative outcome. If only \$2 billion were spent each year on sequestration, the chance of resolution by 2025 would be 58 percent. On the other hand, spending the full \$4 billion on renewables through both periods would lead to a 25 percent chance of no advance, a 50 percent chance of moderate advance, and a 25 percent chance of significant advance. If only \$2 billion were spent each year, there would be a 42 percent chance of no advance, a 45 percent chance of moderate advance, and 13 percent chance of significant advance.

In this example, the decision model would consider whether the combination of the two halfway scenarios outperforms either unilateral strategy. In general, adjusting the allocation across both programs and periods, the model selects the investment strategy that maximizes expected utility, a function of the distribution over the range of possible technological outcomes. The basis of this valuation is discussed next.

6.2 The Outcome Scenarios

Three technological scenarios are possible after the twenty years of R&D, and each scenario aggregates multiple research outcomes, that is, patterns of success and failure between the two research programs. The contraction of the possible outcome space into a discrete set is attractive because it simplifies analysis without necessarily over-simplifying the assumptions about the economic and technological impacts of innovation. The converse of this benefit is that finer distinctions are lost by the aggregation. In particular,

[17] \$4 billion is entirely arbitrary and was chosen for illustrative purposes only.

while themselves intermediate and part of a long term path, the scenarios are in the short term path independent; that is, the particular pattern of success and failure that produced the outcome does not influence its utility valuation. This is a natural area of extension to the model, since it is likely that some patterns are preferable to others.[18] Table 1 illustrates the aggregation, and the following paragraphs describe the scenarios, highlighting as well some of their limitations.

First Period Research Outcome		Second Period Research Outcome		
Sequestration	Renewables	Sequestration	Renewables	Outcome
Positive	N/A	N/A	N/A	First Best
Negative	Success	N/A	Success	First Best
Negative	Success	N/A	Failure	Moderate
Negative	Failure	N/A	Success	Moderate
Negative	Failure	N/A	Failure	Worst Case
Failure	Success	Positive	N/A	First Best
Failure	Success	Neg/Fail	Success	First Best
Failure	Success	Neg/Fail	Failure	Moderate
Failure	Failure	Positive	N/A	First Best
Failure	Failure	Neg/Fail	Success	Moderate
Failure	Failure	Neg/Fail	Failure	Worst Case

Table 1: The Aggregation of Outcomes

First Best: In the first-best scenario, there exists a proven, economically attractive option for reducing GHG emissions. The first-best outcome might arise either by proof of the permanency of carbon sequestration, allowing the current fossil fuel infrastructure to remain largely intact, or by significant advances in renewables, allowing either a slow or, if necessary, a quick transition to these zero-GHG technologies. There is no accounting of transition time, nor of the development of supplementary technologies associated with the transition.[19] Also, this scenario assumes no additional benefit to having

[18]For example, earlier successes, apart from their ability to inform future decisions via updated expectations (which is captured by the model), may be more beneficial because they facilitate earlier deployment and learning.

[19]For example, storage or demand-response technologies may be required to reduce the impacts of renewable intermittency.

both a positive sequestration resolution and significant renewable advance.

Moderate: This scenario is characterized by moderate advances in renewables accompanied by either a failure to resolve the uncertainty regarding the permanency of sequestration or a negative resolution. Renewables can be an important contributor to GHG mitigation, but the costs will not be trivial, and sequestration is not a clear-cut option either because it is highly uncertain or is known to have limited permanency. These are not distinguished by the moderate scenario, but it is likely that the latter case is preferable since it enhances the state of knowledge. Also, the timing issue is not addressed—renewable success in the first period followed by failure is effectively equivalent to failure followed by success.

Worst Case: This scenario is characterized by the lack of an attractive mitigation option. There has been only limited advances in renewables, leading not only to marginal deployment, but also to a sense that additional advances may not be forthcoming. Further, either there remains large uncertainty regarding the permanency of sequestration, making it a risky option, or it has proven to be technically infeasible. In other words, the current state of technology remains largely unchanged by the end of the second period, although our beliefs about the potential for advance in renewables may be different. The same caveat may be made about differentiating between a failure to resolve the sequestration uncertainty and a negative resolution; the timing of such a resolution may also affect the scenario's valuation.

The valuation of each scenario depends on more than its technological characterization. Also affecting society's relative preferences is the extent to which damages from climate change are realized, or are believed to be realized. At present, there is great uncertainty about the impacts of continued GHG emissions on human welfare. These impacts are therefore included as an exogenous uncertainty that affects the valuation of the outcomes. Climate damages, or more precisely, our beliefs about climate damages by the end of the time horizon, may resolve to one of three states, each of which implies a different utility weighting for the three outcomes. First, damages may remain in the predicted range, characterized by a relatively tight clustering of utility values. Second, damages may appear to be severe, which primarily affects the worst-case technological scenario, because severe damages, or the stabilization effort required to avoid them, will be extremely costly if little advance has been made in renewables and sequestration has proven technically infeasible. Third, there may be no climate damages, so that their effect

Technological Scenarios	Climate Damages		
	Predicted	Severe	Neutral
Probability	0.6	0.3	0.1
First Best	1	1	1
Moderate	0.5	-1	1
Worst Case	0	-10	1

Table 2: Relative Utility Valuation of Technological Outcomes with Climate Uncertainty

on the technological outcomes is neutralized, implying an equal weighting of all three. Climate risk is introduced by applying a probability distribution to these three resolutions. In contrast, a control case is considered where only the predicted damages are possible.

Table 2 shows the utility values used in the analysis. These values are meant merely as placeholders to allow a demonstration the model. As discussed above, IAMs and other modes of evaluations could be applied in future revision and expansion of the model to establish a more robust quantification of the utility matrix. For this framework to be ultimately viable, it is critical that the utility valuation employed reflect a relative preference set that is not only defensible, but also transparent.[20] A perhaps more tractable representation of utility values can be made by considering hypothetical lotteries. For example, the numerical values in Table 2 imply that under severe climate damages, society would be indifferent between obtaining the moderate outcome and an uncertain lottery with an 82 percent chance of obtaining the first-best outcome and an 18 percent chance of obtaining the worst-case outcome. Depending on the definition of "severe climate damages", this perspective may be a reasonable starting point from which to consider the values of the different technological outcomes.

The values used in the analysis here also imply that in the first-best technological scenario, the extent of climate damages is economically irrelevant. That is, increased climate sensitivity, with correspondingly lower concentration targets, does not increase the cost of emissions reductions, because either

[20]Note that the absolute values of the utility weighting are not the appropriate index for interpretation. Since the weighting schemes represent a utility function, they need only be consistent among themselves. Any linear transformation of the weights will provide the same result.

sequestration or renewables acts as a backstop technology providing essentially limitless supply at constant marginal cost. In the case of no significant climate damages, the utility values imply that all technological scenarios are equivalent, because there will be no need to take any action. However, it is possible that society should value this set of outcomes strictly less than the set defined by the first-best technology scenario, since the investment in R&D has proved futile. Such opportunity costs are not considered elsewhere in the model because of the assumption that the budget to be allocated is fixed.

7 Results

This section discusses the results of analysis conducted using the stylized application. We begin by exploring, in broad terms, the circumstances that would lead to a specialized optimal portfolio—one weighted heavily toward one of the two research programs. Comparative statics are explored next, in particular for the parameters ρ (referred to as productivity, since it scales the success probability) and q (conditional probability of a positive sequestration uncertainty resolution).[21] Following the questions posed in the Introduction, we first discuss the forces that drive the relative allocation between the two programs, and second address the impacts of uncertainty. A discussion of the dynamic response of optimal R&D to research success or failure follows, and we end with a brief discussion of results from the numerical example.

As pointed out earlier, no strong conclusions can be drawn about the precise characteristics of an optimal allocation. Rather the results suggest some possible dynamics between the two research programs under consideration, along with more general results about diversification that can be inferred from the pairwise comparison. Overall, the main value of the results presented here is to demonstrate the types of insights available from this conceptual approach, and to show what steps may be taken to produce a more rigorous set of findings.

[21] Further work should examine the effects of varying other parts of the model, such as the functional form used in the R&D process model and elements of the utility valuation matrix.

7.1 Specialization vs. Diversification

Within the structure of the model, specialization occurs when one or more of the underlying parameters are equal to zero or one; that is, either when one of the research programs is wholly ineffective at achieving its objective, or when the *a priori* belief is that sequestration has no chance for significant permanency. The case of zero productivity or zero chance of a positive resolution is trivial—naturally all investment will go to the other program. Similarly, when $\rho_{rnw} = 1$, the first-best outcome can be reached with probability one by investing completely in renewables in both periods, so that sequestration is not necessary. However, even when $\rho_{seq} = 1$, as long as $q < 1$, there remains uncertainty about a positive resolution, so that some investment in renewables is *always* optimal (provided $\rho_{rnw} > 0$). This is an important point that is driven expressly by the success structure of the sequestration research program, and is therefore robust to most other model assumptions.[22]

For all non-boundary values of the parameters, some diversification is optimal, that is, both programs receive positive investment. This result is driven by the decreasing returns to scale assumption, as discussed above, and is also consistent with the motivation to maintain capabilities in both areas. The variation of the optimal strategy in response to changes in the productivity parameters in several settings are pictured in a series of figures.

7.2 Near-term Renewables vs. Sequestration

A first question concerning the pairwise comparison is that of the postponement of renewable energy technology development until the retention rates of geologically sequestered carbon are better understood. As discussed above, the model's results suggest in the realistic case that neither program is guaranteed to succeed, some investment in renewables is always optimal to balance the portfolio and provide insurance against the failure of sequestration. Thus the model does not support the case for postponement, at least not in the strict sense of complete cessation of renewables R&D. This effect is even more pronounced when uncertainty about climate damages is introduced, as discussed below.

[22] Specialization could also arise if R&D success had negligible value—for example, if even two successes in renewables led only to the worst-case outcome. The stylized model discussed here is structured so that this basis for specialization cannot occur, although the case where $q = 0$ can be seen in this light.

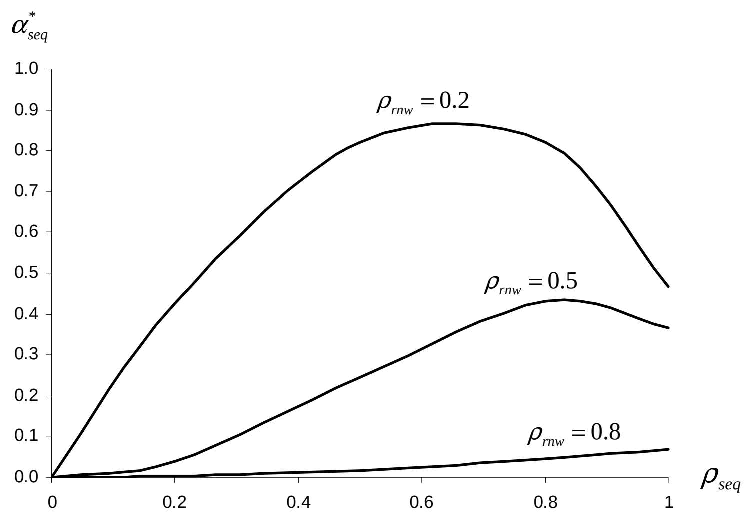

Figure 3: First Period Investment With No Climate Uncertainty

A second question is related and asks specifically whether optimal near-term investment in resolving the sequestration uncertainty should increase in beliefs about its productivity or chance of a positive resolution. Unsurprisingly, sequestration investment is always increasing in the latter belief, that is, in the parameter q. In most cases, the same is true for the productivity parameter ρ_{seq}. Given that an optimal strategy is being pursued, if new information suggests that either of these parameters is greater than previously thought, the new optimum will include a greater investment in the sequestration program. This relationship is reflected by the positive slope of the curves in Figure 3. However, when sequestration productivity is believed to be very high, with productivity in renewables considerably less so, the optimal allocation to sequestration actually declines in the productivity of sequestration research. In other words, the optimal response to the same new information as above would in this case be to invest *less* in the sequestration program. The result may be seen graphically in the downturn of the upper two curves in Figure 3.

This is perhaps the most striking result encountered by the analysis. One intuitive explanation is that when confidence in the success of a sequestration program is high, the insurance provided by the renewable program more than offsets the opportunity cost of waiting until the second period. In most instances of this phenomenon, optimal second period investment in sequestration is close to 100 percent, suggesting that the deferral option drives the result. However, at this point, we do not consider the dynamic well understood, and recommend it as a topic for continued research.

7.3 Uncertainty about Climate Damages

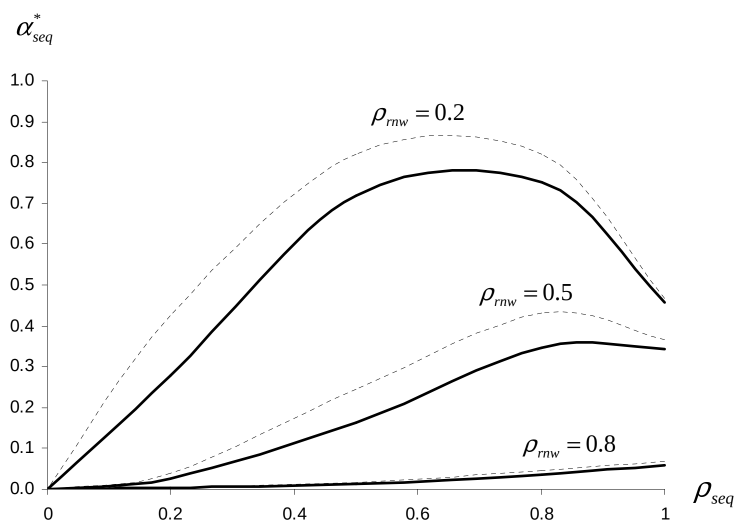

Figure 4: First Period Investment With Climate Uncertainty

Another important question is the response of the optimal investment mix to increasing uncertainty about the appropriate stabilization target, or climate damages in general. To examine this question, the probability distribution shown in Table 2 was applied to the climate damage outcones and compared to a control case were only predicted damages were possible. With the introduction of uncertainty about climate damages, investment in sequestration is depressed in both periods. Figure 4 demonstrates this effect for first period decision, with the control case represented by a dotted line and the uncertain case represented by the solid line.

This result is driven primarily by the dramatic decrease in utility in the worst-case scenarios under severe climate damages (recall the hypothetical lottery illustration), and by the fact that investment in the renewable program, as opposed to the sequestration program, allows for attainment of the moderate scenario. To the extent that these two structural attributes are realistic, the model provides a powerful insight into the relationship between R&D planning and climate uncertainty. Moreover, these results have implications on the response of the optimal investment strategy to other kinds of exogenous uncertainty, highlighting the role of risk management in motivating diversification. In this case, increased risk aversion (or equivalently, increased risk with constant risk aversion) suggests an optimal diversification strategy of more heavily favoring research programs with the possibility of moderate outcomes.

7.4 Dynamic Response

A question not inherently related to these two specific research programs, but interesting nonetheless for R&D planning in general concerns the implications of first period performance for the second period investment decision. Some intertemporal dynamics have been suggested in this section, but we have not directly evaluated the claims made in Section 3 about the benefits of information about technological potential for future decisions. In fact these benefits are clearly demonstrated by the results for the entire range of the model's parameters—the optimal investment strategy in the second period is significantly altered by whether research programs succeed or fail in the first period, implying that such information is useful in optimizing the dynamic response.

Figures 5 and 6 show the optimal allocation decision in the second period when the the renewables program has succeeded and failed in the first period, respectively. In the case of first period success in renewables, optimal investment shifts dramatically toward that program in the second period. Note that when the carbon sequestration program succeeds, the second period investment decision, as it is stated in this model, disappears, because there is no longer a need for uncertainty resolution. Thus these intuitive results suggest that early successes in research programs, without actually changing beliefs about the future productivity (the ρ parameters remain constant from period to period regardless of research outcomes), can justify increased investment through the information they provide about what can be achieved by the end of the time horizon.

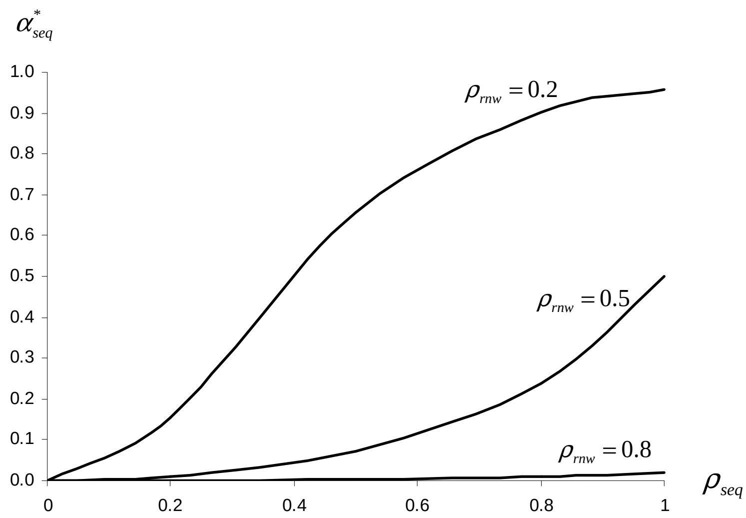

Figure 5: Second Period Investment after First Period Success in Renewables

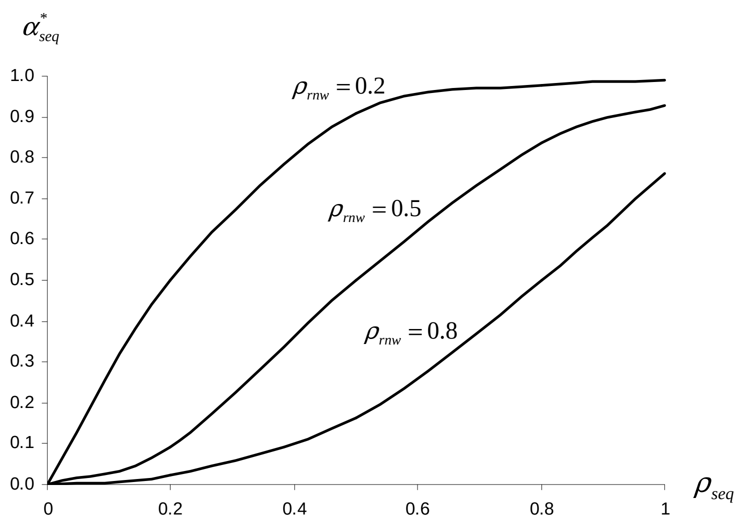

Figure 6: Second Period Investment after First Period Failure in Renewables

7.5 Numerical Example

To conclude this discussion, the model was applied to the numerical example in Section 6.1 with all three parameters equal to 0.5 (and no climate uncertainty). Under these circumstances, the optimal near-term decision is to weight the portfolio most heavily towards renewables, with less than $1 billion invested annually to resolve the permanency of sequestration. In the second period, if the renewable program has failed, the optimal strategy is perfect balance, while after a success, the renewable program receives more than 90% of optimal investment. Fundamentally, however, the optimal allocation of investment is too dependent on the an accurate characterization of the productivity of R&D and the weighting of the outcomes to be satisfactorily calculated here. The goal of this paper has been to present a framework in which these fundamental factors can be better understood.

8 Conclusion

There are several significant conclusions that can be drawn from the work described here. First and foremost, there is a research need for a better understanding of the R&D process as it relates to the climate change context. The existing literature has not yet broached the subject, and large investments are currently being made, with even larger expenditure required in the future, in the development of new energy technologies. Despite the complexity and subjectivity inherent in the problem, these investment decisions can be informed by R&D planning models.

The approach here is an important first step in demonstrating an analytical framing of the problem. To a certain extent it is a canonical problem, with a structure independent of the climate context, although few other instances of similar magnitude and breadth exist today. The key aspects of the problem have been identified: describing research programs in terms of the purpose of the research, the process by which investment produces results, and the valuation of those results. Stylized analytical mechanisms have been developed for each model component that capture the main drivers while remaining tractable. In particular, the scenario mechanism provides an attractive connection to other analysis in the field, such as that conducted with large-scale IAMs. Finally, sketching only the details of these mechanisms allowed the extraction of several interesting insights about the nature of the optimal allocation decision across time.

An important theme that emerges strongly from this exercise is the need for more specific information on the potential for technological advance, and on its relationship to R&D effort. Ironically, the most compelling conclusion of this study may be the inability to make any robust conclusions about an optimal R&D strategy with the current state of knowledge about the R&D process. Parallel to the development of a deeper understanding of technologies themselves, continued refinement of the analytic framework presented here should focus on effective incorporation of this technological information.

The application in this paper is a prototype. Both the R&D process model and the outcome scenarios require revision and expansion. To this end, the key next steps include: (1) more rigorous, defensible, and transparent valuation of outcomes based on, for example, linkage to IAMs or expert solicitation; (2) the development of a more accurate innovation production function, perhaps with an empirical basis; (3) improving the resolution of the outcome scenarios to include timing and path dependence issues; (4) exploring R&D programs with variable time horizons; (5) accounting for technological outcomes exogenous to the R&D process, such as enhanced public acceptance of nuclear energy; and (6) considering the optimal size of the R&D budget as well as the relative allocation. It is our hope that, apart from these developments, the work presented here may serve as departure point for a wider analytic consideration of the R&D planning problem, an instrumental part of the climate challenge.

References

Baker, E., L. Clarke, and J. Weyant (2003). R&D as a hedge against climate damages. *Submitted to Energy Journal*.

Caldeira, K., A. Jain, and M. Hoffert (2003). Climate sensitivity uncertainty and the need for energy without CO_2 emission. *Science 299*, 2052–2054.

Dasgupta, P. and E. Maskin (1987). The simple economics of research portfolios. *The Economic Journal 97*, 581–595.

Edmonds, J., M. Wise, and C. MacCracken (1994). Advanced energy technologies and climate change: n analysis using the global change assessment model (gcam). Technical Report PNL-9789, UC-402, Pacific Northwest National Laboratory, Richland, WA.

Goulder, L. and K. Mathai (2000). Optimal CO_2 abatement in the presence of induced technological change. *Journal of Environmental Economics and Management 39*(1), 1–38.

Goulder, L. and S. Schneider (1999). Induced technological change and the attractiveness of CO_2 abatement policies. *Resource and Energy Economics 21*, 211–253.

Gritsevskii, A. and N. Nakicenovic (2000). Modeling uncertainty of induced technological change. *Energy Policy 28*.

Grubler, A. and A. Gritsevskii (2001). A model of endogenous technological change through uncertain returns to learning (R&D and investments). *Ukranian Economic Review*.

Hoffert, M., K. Caldeira, G. Benford, D. Criswell, C. Green, H. Herzog, A. Jain, H. Kheshgi, K. Lackner, J. Lewis, H. Lightfoot, W. Mannheimer, J. Mankins, M. Manuel, L. Perkins, M. Schlesinger, T. Volk, and T. Wigley (2002). Advanced technology paths to global climate stability: Energy for a greenhouse planet. *Science 298*, 981–987.

Hoffert, M., K. Caldeira, A. Jain, E. Haites, L. Harvey, S. Potter, M. Schlesinger, S. Schneider, R. Watts, T. Wigley, and D. Wuebbles (1998). Energy implications of future stabilization of atmospheric CO_2 content. *Nature 395*, 881–884.

IPCC (2001). *Climate Change 2001: Mitigation*. Cambridge University Press, New York, NY.

Keeney, R. and H. Raiffa (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, New York.

Loch, C. and S. Kavadias (2003). Dynamic portfolio selection of npd programs using marginal returns. *Management Science 48*(10), 1227–1241.

Manne, A., R. Mendelsohn, and R. Richels (1993). Merge: A model for evaluating regional and global effects of ghg reduction policies. In N. Nakicenovic, W. Nordhaus, R. Richels, and F. Toth (Eds.), *Integrative Assessment of Mitigation, Impacts, and Adaptation to Climate Change*. Laxenburg, Austria: International Institute for Applied Systems Analysis (IIASA).

Nordhaus, W. (1994). *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA: MIT Press.

Nordhaus, W. (2002). Modeling induced innovation in climate-change policy. In A. Grubler, N. Nakicenovic, and W. Nordhaus (Eds.), *Technological Change and the Environment*. Washington, D.C.: Resources for the Future.

Popp, D. (2002). ENTICE: Endognenous technological change in the DICE model of global warming. Syracuse University, Submitted to *Journal of Environmental Economics and Management*.

Prinn, R., H. Jacoby, A. Sokolov, C. Wang, X. Xiao, Z. Yang, R. Eckaus, P. Stone, D. Ellerman, J. Melillo, J. Fitzmaurice, D. Kicklighter, G. Holian, and Y. Liu (1999). Integrated global system model for climate policy assessment: Feedbacks and sensitivity studies. *Climatic Change 41*(3/4), 469–546.

Romer, P. (1990). Endogenous technical change. *Journal of Political Economy 98*(5), S71–S102.

Rosenberg, N. (1990). Why do firms do basic research (with their own money). *Research Policy 19*(2), 165–174.

Schneider, S. and L. Goulder (1997). Achieving low-cost emissions targets. *Nature 389*(4), 13–14.

Seebregts, A., T. Kram, G. Schaeffer, A. Stoffer, S. Kypreos, L. Barreto, S. Messner, and L. Schrattenholzer (1999). Endogenous technological change in energy systems models: Synthesis of experience with eris, markal, and message. Technical report, Netherlands Research Foundation ECN.

van der Zwaan, B., R. Gerlagh, G. Klaassen, and L. Schrattenholzer (2002). Endogenous technological change in climate change modeling. *Energy Economics 24*, 1–19.